LA-UR -88-648

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

LA-UP--88-648

DE88 006447

TITLE KNOWLEDGE-BASED SYSTEM FOR COMPUTER SECURITY

AUTHOR(S) William J. Hunteman

SUBMITTED TO 11th National Computer Security Conference, Baltimore, MD, October 17-20, 1988

DISCLAIMER

Los Alamos Los Alamos National Laboratory
Los Alamos, New Mexico 87545

Knowledge-Based System for Computer Security

William Hunteman
Los Alamos National Laboratory
Los Alamos, NM

ABSTRACT

The rapid expansion of computer security information and technology has provided little support for the security officer to identify and implement the safeguards needed to secure a computing system. The Department of Energy Center for Computer Security is developing a knowledge-based computer security system to provide expert knowledge to the security officer. The system is policy-based and incorporates a comprehensive list of system attack scenarios and safeguards that implement the required policy while defending against the attacks.

INTRODUCTION

The field of computer security has undergone a significant expansion of the information and technology available to address security concerns in computing systems. The advances are often directed towards technological solutions of a multidimensional problem, but the nontechnical areas of the computer security field have received little, if any, serious effort towards improving the quality of security. This paper describes an effort under way at the Department of Energy (DOE) Center for Computer Security to create a knowledge-based system to act as an advisor to the security officer when developing a secure system or reviewing security in an existing system.

PROBLEM

The security officer is required to integrate the advances in computing system hardware and software with the users' need for productivity while attempting to implement security policies that are at best complex, difficult to interpret into the local environment, and often apparently in conflict with each other. The security officer in today's computer security environment is faced with a large quantity of information in the form of written and unwritten policies, large software systems, and increasingly complex hardware. This information is frequently uncertain or incomplete, and the security officer is required to interpret it into a set of safeguards for a particular computer system with the goal of "securing" the system against some poorly specified threats.

The lack of meaningful education and assistance for the average security officer frequently requires the individual to resort to reliance on tradition or folklore to secure a system. Many security officers have been assigned previously "secured" systems and are expected to continue the tradition with little or no understanding of what was done or why it was done when the system was originally secured. Frequently, when a system is expanded or altered the first security officer activity is to talk with other security officers to

seek advice about how similar systems were secured. The goal seems to be "tell me what is the minimum I must do to get my system accredited."

Security officers who have access to computer security knowledge because of their own experience or through experts are able to implement and maintain better, more comprehensive security programs. The primary need is to collect, organize, and present the security knowledge of the experts. The information must be presented in a manner tailored to the security officer's requirements, and the officer must be able to query the "expert" for a justification or explanation of a decision or recommendation.

Development of a knowledge-based system incorporating a methodology to combine uncertain or incomplete information and manage the large quantities of "expert" knowledge will extend the information to any security officer on demand.

DESIGN GOALS

The Knowledge-Based System for Computer Security is designed to provide an integrated collection of policy requirements and expert knowledge to the security officer. The design goals for the system are:

- Define a comprehensive list of safeguards based on policy requirements.
- Collect a detailed description of the local computing environment.
- Produce a list of safeguards that are applicable to the local computing environment with guidance on the required implementation approach for each safeguard.
- Effectively address the problem of uncertain or incomplete information in the knowledge bases.
- Support "what-if" experimentation to allow the security officer to adjust the local computing environment or reject specific safeguards because of resource limitations.
- Provide, on request, justification or explanation of each decision throughout the process.
- Provide a user interface oriented to the security officer's needs and environment.

SYSTEM ARCHITECTURE

The system architecture (Figure 1) is composed of several knowledge bases and an inference engine designed to access the knowledge bases and to provide an appropriate interface for the security officer.

Each of the knowledge bases - Policy, Data Exposure, Sensitivity Level, Attack, and Computing Environment - reside in independent data files that are accessed by the inference engine. The use of external data files allows rapid updating of the knowledge to reflect policy changes, technology advances, or adjustments in the attack scenarios.

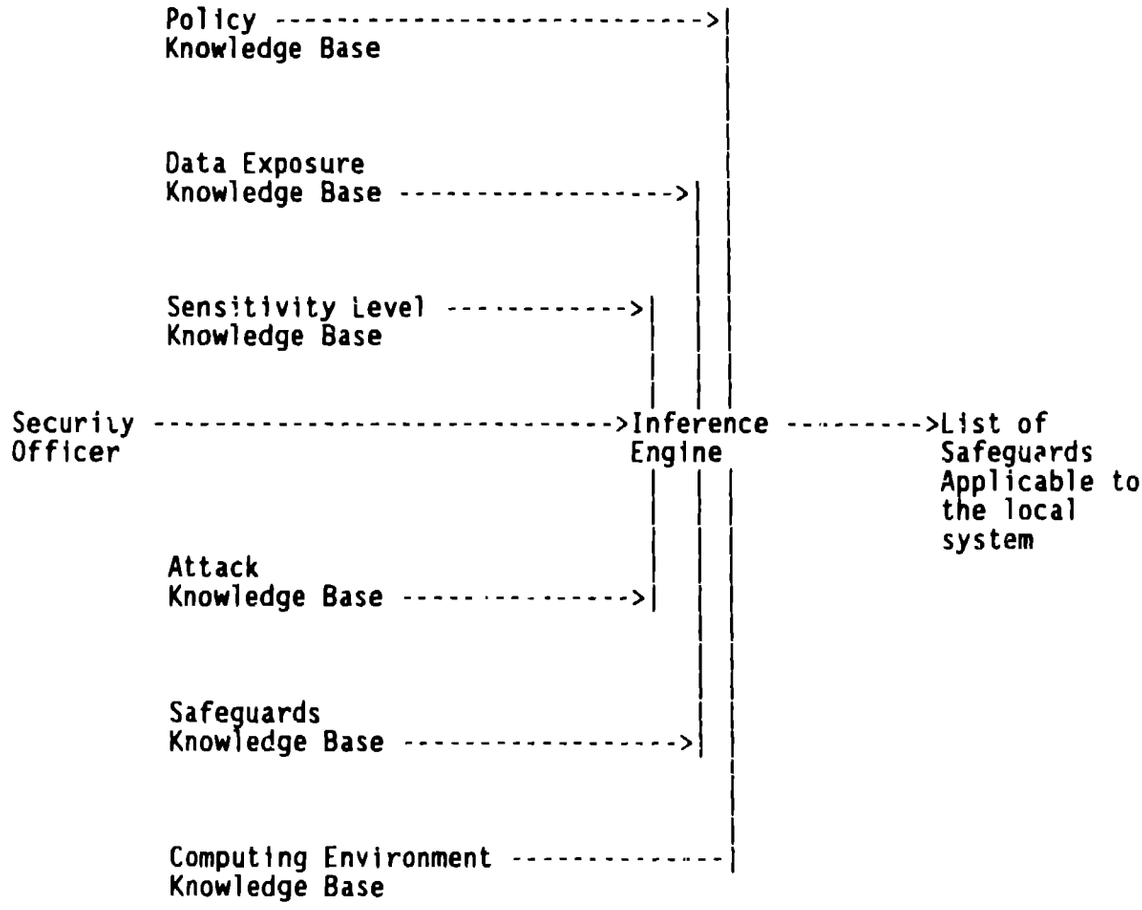

Figure 1.

System Architecture

POLICY KNOWLEDGE BASE

Most policy statements are complex and difficult to interpret for a local environment. Experts from the policy-making organization will often give conflicting advice regarding how the policy should be implemented in a particular computer system. The lack of clear guidance on applying the policies and the inconsistent implementations clearly suggest that a uniform methodology for interpreting and applying a policy is needed. The methodology and system being developed provide a consistent decomposition of policy statements into a knowledge base that is used to develop guidance for implementing safeguards in a specific computing system.

The Policy knowledge base (Figure 2) contains specific security requirements obtained by decomposing information from all applicable policy statements and other sources, e.g. inspection or evaluation criteria and good practices. When appropriate, the Policy knowledge base also can be extended by local regulations or standards that augment the national policy requirements.

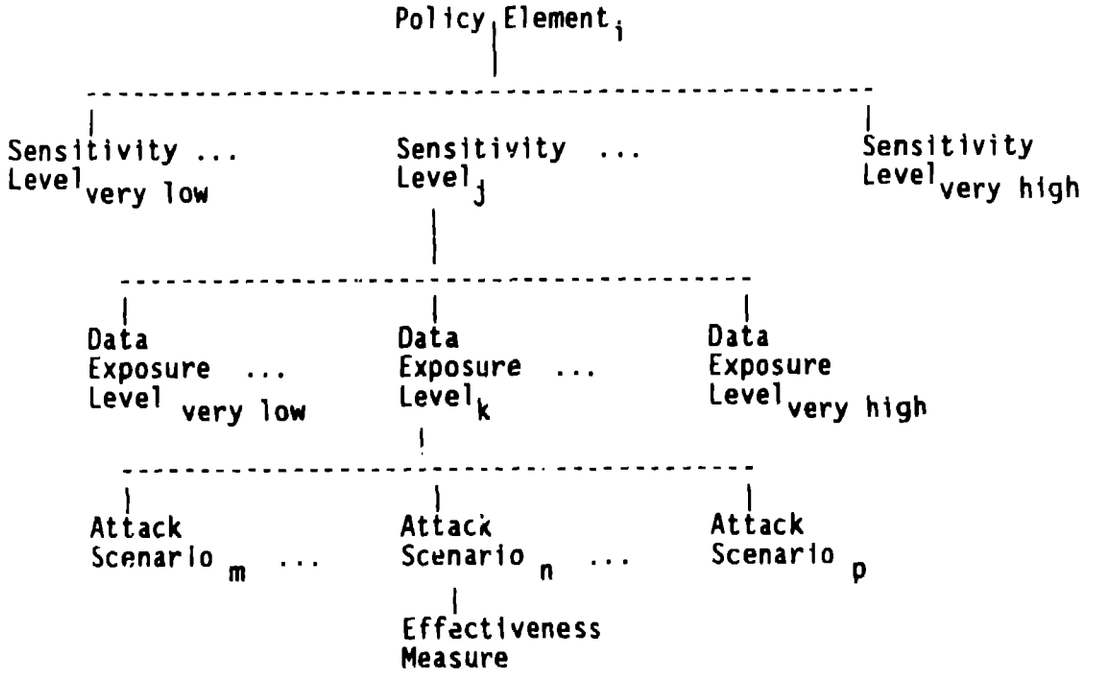

Figure 2.

Policy Knowledge Base

Each policy element is a single policy requirement, e.g., each user of an ADP system must be authenticated before access is permitted. Each policy element is evaluated against all attack scenarios in the Attack knowledge base. Each evaluation includes all sensitivity and data exposure levels. The evaluation is based on expert assessment, the establishment of confidence factors, and the use of established techniques for combining uncertain or incomplete evidence.

The evaluation activity results in a measure of effectiveness for each policy element, sensitivity level, data exposure level, and attack scenario tuple.

Figure 3 depicts an example of a policy element entry in Policy knowledge base where:

Policy element = User must be identified and authenticated before access to the system is permitted.
Attack scenario = Disclosure of information
Failure of access controls
Unauthorized access to information
Impersonation of authorized user

This example is not complete and is used only to illustrate the concepts involved in the Policy knowledge base.

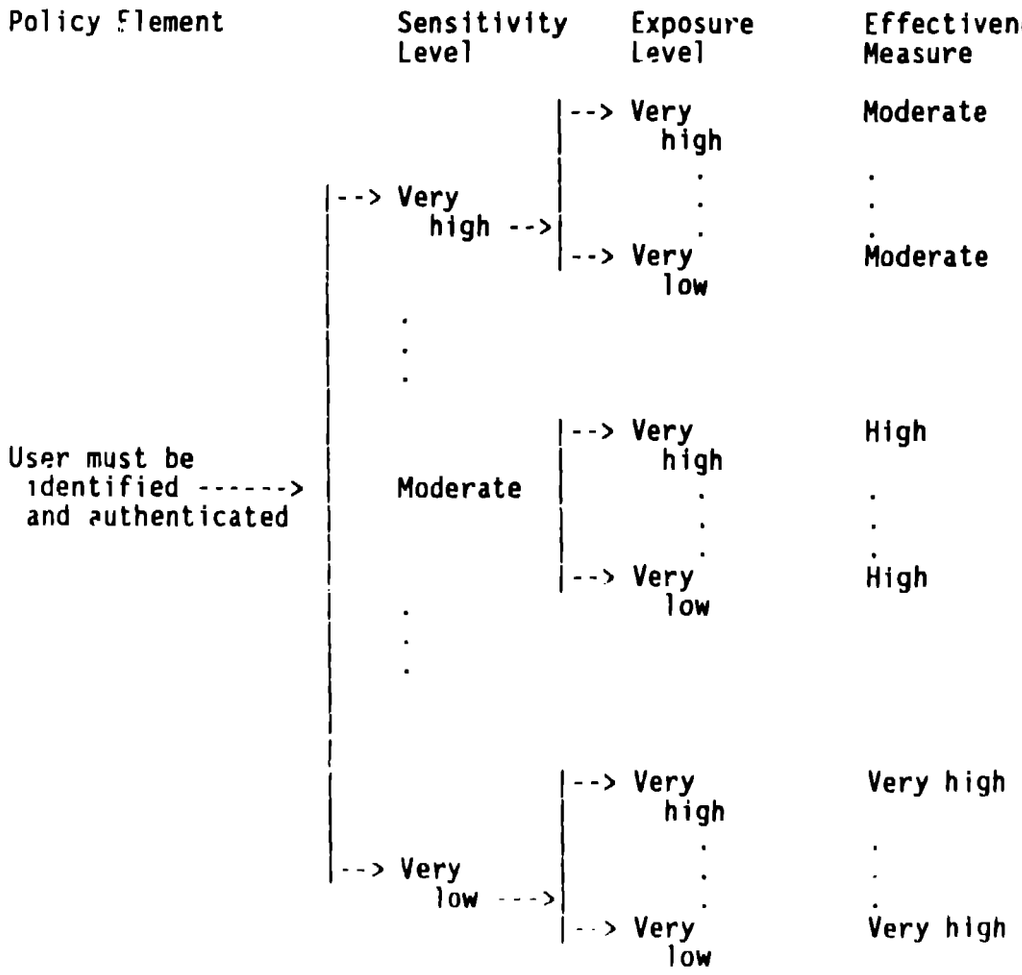

Figure 3.

Policy Knowledge Base Example

DATA EXPOSURE KNOWLEDGE BASE

Many computer security policies either fail to, or only partially, address an assessment of the degree of exposure of information in a computer system. Determination of the degree of exposure provides an important criterion for use in evaluating the required safeguards and the level of implementation needed for each safeguard. The Data Exposure knowledge base defined for this system is designed to provide a measure of the potential exposure of information stored or processed on the computing system.

The Data Exposure knowledge base contains possible relative data exposure levels, ranging from very low to very high. These are used to assess the effectiveness of policy elements and safeguards against the various attack scenarios and to develop a measure of the data exposure in the local computing environment.

SENSITIVITY LEVEL KNOWLEDGE BASE

Every security officer is required to assess the sensitivity of the information stored or processed on the system during the process of securing the system. Often this determination is automatically derived from the classification and category of information on the system, but other local factors, such as local or national politics, or the cost of collecting the information, may dictate a higher level of sensitivity than the data classification may suggest. The sensitivity measure is another element that the security officer must include when determining what safeguards must exist and the level of implementation for each safeguard.

The Sensitivity Level knowledge base contains a range of possible relative sensitivity levels that are used to assess the effectiveness of the policy elements and safeguards against the various attack scenarios and to develop a measure of sensitivity in the local computing environment.

ATTACK KNOWLEDGE BASE

The Attack knowledge base contains information covering all forms of possible attack scenarios. The attack scenarios are developed in consultation with experts using a top-down approach beginning with the traditional concerns of disclosure, destruction, distortion (modification), delay, and denial of use. Figure 4 depicts a portion of the attack scenarios and the refinement approach used to define the various scenarios.

The attack scenario approach was chosen because traditional threat-based perspectives have several serious problems. Frequently, threat information with sufficient detail to assist the security officer in selecting safeguards for the local system is not available due to the lack of sufficient clearances by the security officer or the need to protect intelligence sources. Often the threat guidance produced by the policy organizations is either sanitized to the point of having little value, or is simply a recitation of the identified specific threat sources with no guidance on how to interpret the information for a local computing system.

Another difficulty with threat statements is completeness. Often a threat statement is interpreted by a security officer as the "recipe" of concerns to be addressed. If a particular attack scenario is not identified in the threat list then it may not be considered by the security officer.

The "recipe" approach also may cause the security officer to focus on specific "threats" while ignoring other large classes of potential system attacks. These classes include such attack scenarios as the public claim by an individual of a successful attack on a particular system. The resources needed to respond to a claimed attack can often match the resources used in responding to an actual attack.

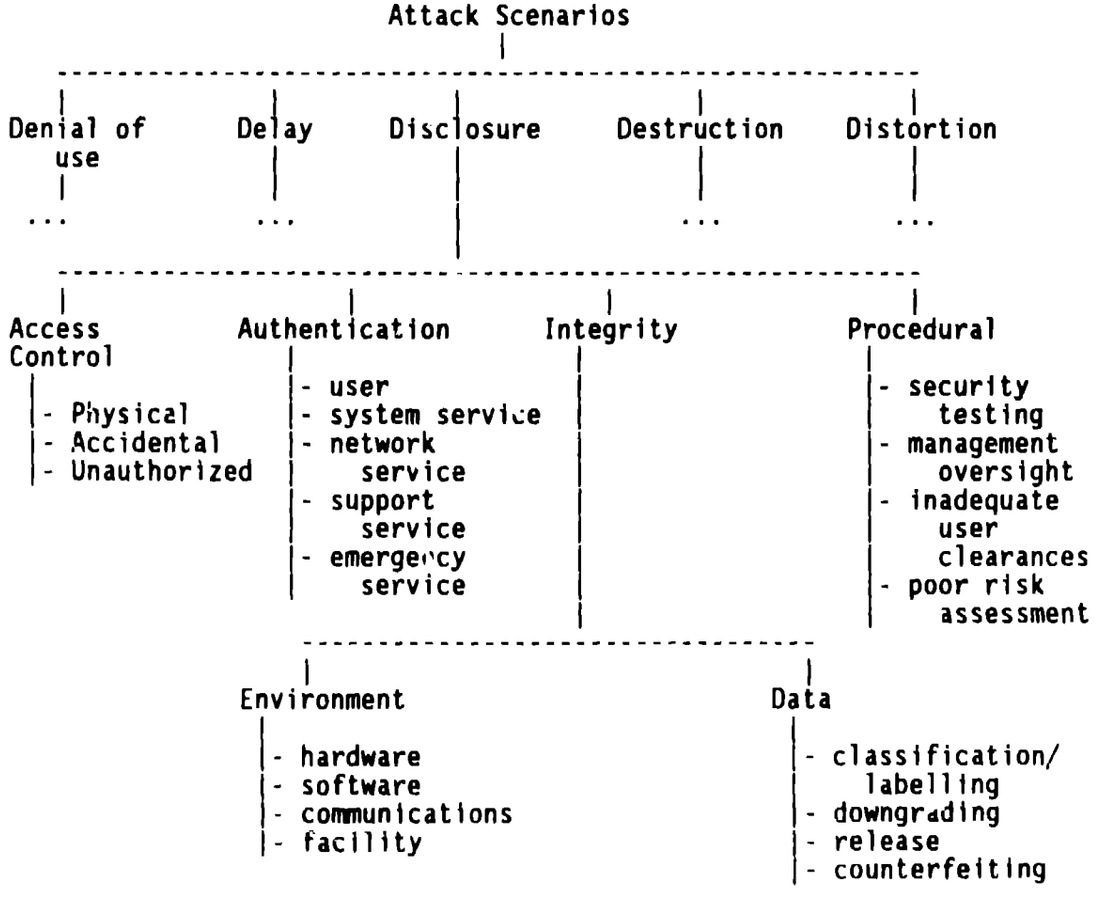

Figure 4.

Attack Scenarios

Each element in the top layer (disclosure, etc.) is decomposed into the four categories (access control, authentication, integrity, and procedures). Each of the categories is further expanded to the subcategories shown in Figure 4.

The subcategories are expanded into more specific attack scenarios. For example, the Access Control category under Disclosure would contain the subcategories Physical, Accidental, and Unauthorized. The Unauthorized subcategory would be further expanded to On-line and Off-line. Attack secnarios that might be included in the On-line area are - mislabelling of information, covert channels, scavenging, wiretaps, viruses, impersonation of authorized user, etc. Attack scenarios that might be included in the Off-line area are - exploiting acts of nature, improper disposal of information in trash, blackmail or coercion, emanations, etc. Each of these attack scenarios would be further expanded to include non-system specific details of the particular form of attack.

SAFEGUARDS KNOWLEDGE BASE

The Safeguards knowledge base (Figure 5) contains a list of potential safeguards derived from computer security experts and existing policies and regulations. Each safeguard is expressed in terms of implementation requirements for each possible data exposure level within each possible sensitivity level. The implementation requirements are defined in terms of the minimum acceptable implementation, an acceptable implementation, and the maximum effective implementation.

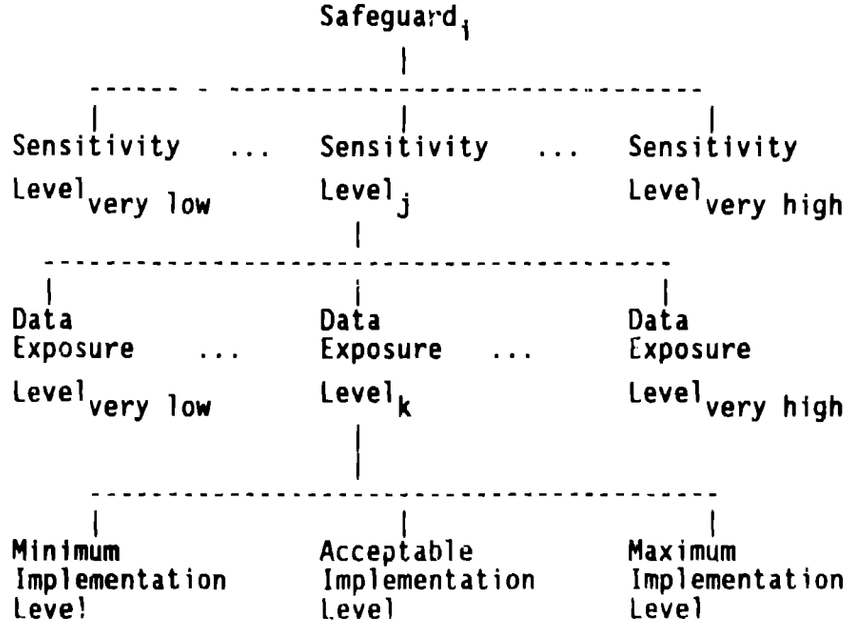

Figure 5.

Safeguards Knowledge Base

The minimum, acceptable, and maximum implementation requirements depend upon the data exposure and sensitivity levels. Each of the implementation levels

is described in a generic fashion relative to the exposure and sensitivity levels. For example, a minimum password implementation for a low exposure level and a low sensitivity level might be defined as six characters in length, machine-generated, and changed once every twelve months. A minimum password implementation for a high sensitivity level might be ten characters in length, machine-generated, and changed every three months.

Many of the implementation elements in this knowledge base will be empty, especially for high data exposure and high sensitivity levels. Very high data exposure and sensitivity levels typically require a sufficiently strong implementation that effectively reduces the options to a single "minimum" acceptable implementation.

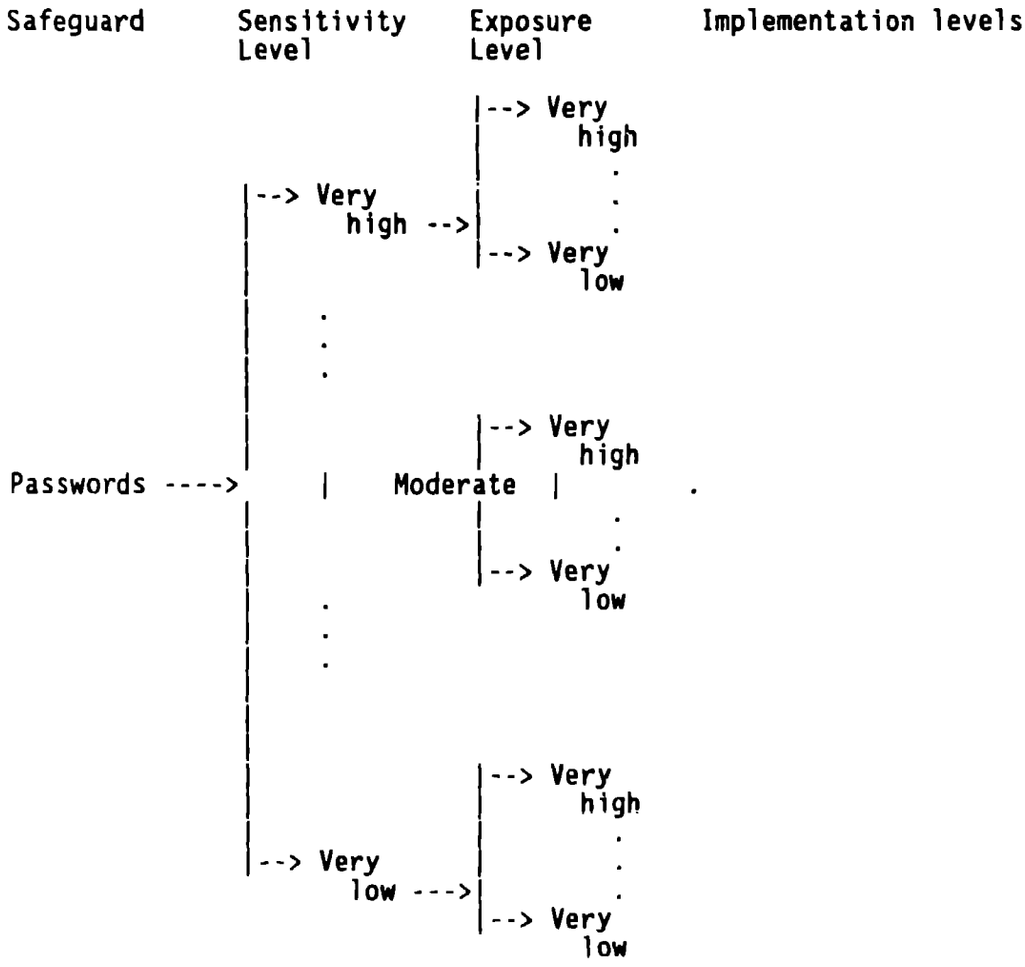

Figure 6.

Example of Information in a Safeguard Knowledge Base Entry.

Once the safeguards and the implementation levels have been identified each safeguard is evaluated against all other safeguards in the knowledge base to determine if the safeguard is equivalent to another safeguard, supports another safeguard, or conflicts w th another safeguard. The evaluation process compares safeguard implementations at a given sensitivity and data exposure level with other safeguards at the same sensitivity and data exposure levels. Equivalent safeguards may be substituted for each other and are important for cost-benefit considerations when a security officer selects the appropriate set of safeguards to secure the computing system. Some safeguards complement others and thereby strengthen the safeguard against the attack scenarios or permit it to be used at a higher data exposure or sensitivity level. Some safeguards may be redundant or conflict with other acceptable safeguards and either reduce or eliminate any security benefit when the safeguards are used together.

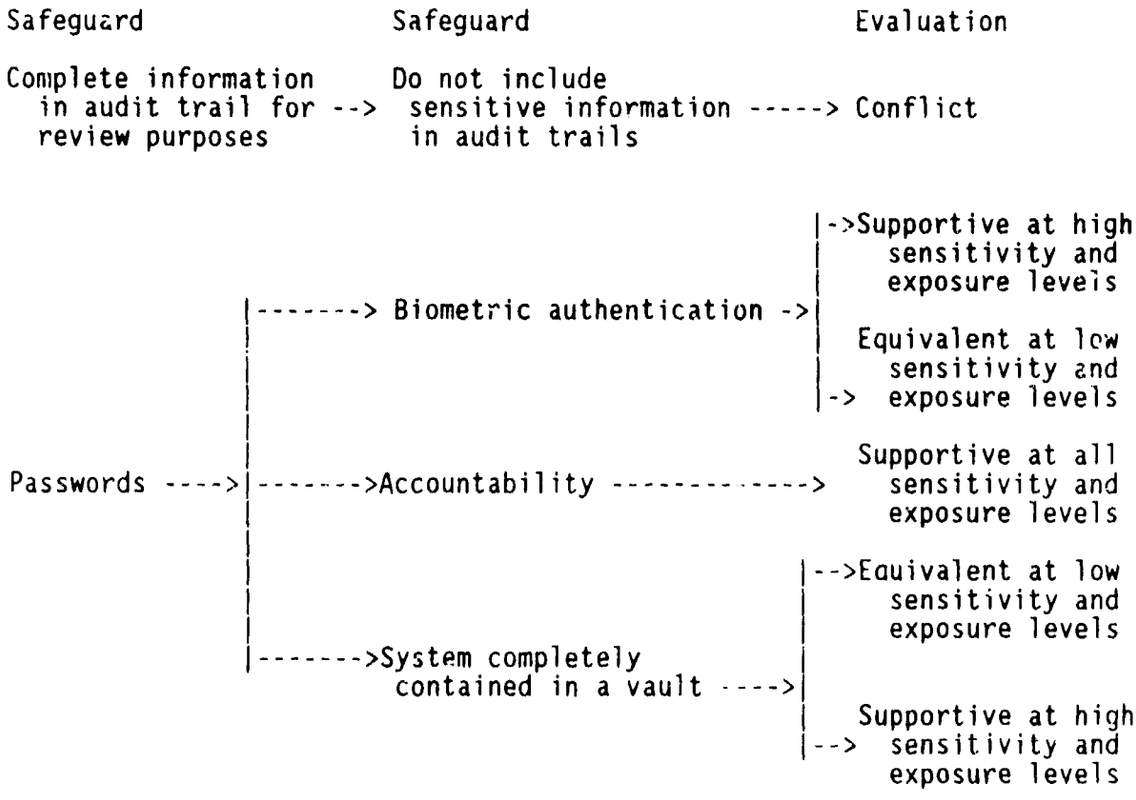

Figure 7.

An Example of Safeguard Comparison with Other Safeguards.

COMPUTING ENVIRONMENT KNOWLEDGE BASE

The Computing Environment knowledge base (Figure 8) is composed of generic descriptions of potential data exposures, sensitivity concerns, and component level definitions of computer systems.

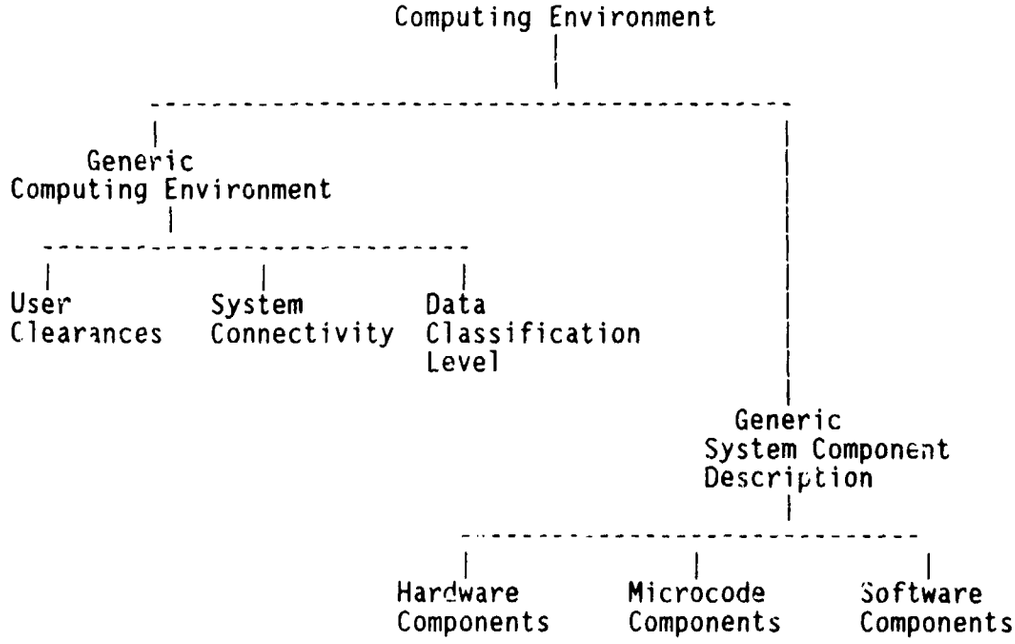

Figure 8.

Computing Environment Knowledge Base

The generic data exposure and sensitivity concerns are used in conjunction with the Data Exposure, Sensitivity, and Attack knowledge bases to collect information from the security officer about the local computing environment. From the security officer responses the system develops a measure of the local data exposure and sensitivity. The Generic Computing Environment knowledge base includes information, such as the range of maximum and minimum user clearances, types of system connectivity (e.g., connections to remote terminals or other computer systems), range of operating modes of the system (e.g., dedicated, system high or multilevel), and the possible highest levels of classification and most restrictive categories of data that might reside on the system.

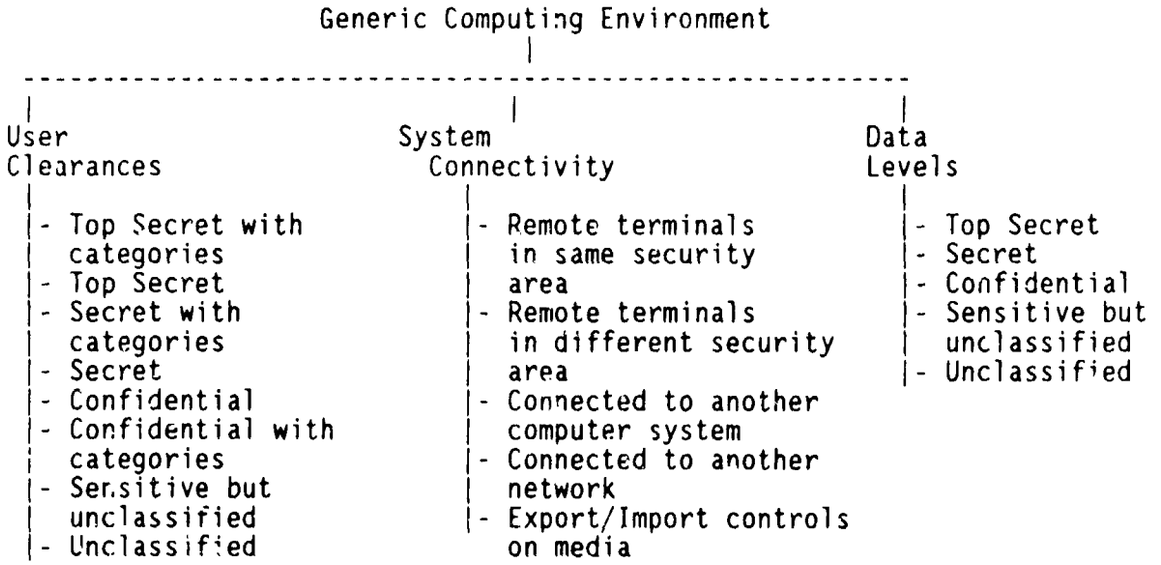

Figure 9.

Example of Information in the Generic Computing Environment Knowledge Base.

The generic description of computer system components is used to formulate questions to the security officer to develop a description of the components in the local computer system. Once the local description is obtained the information is combined with the Attack knowledge base to develop a local level of security concern for each system component. The Generic System Component description knowledge base contains high-level descriptions of the hardware, microcode, and software components found in a typical computing system.

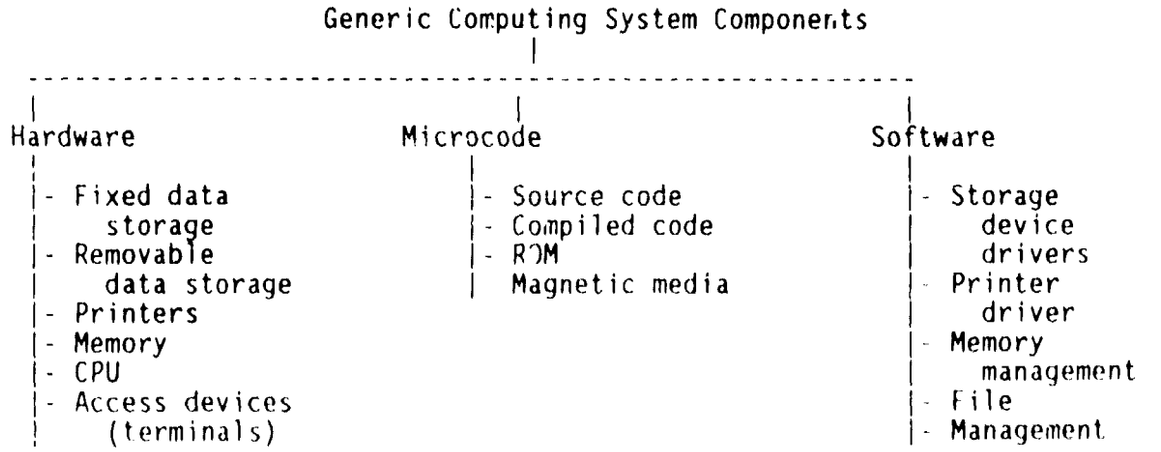

Figure 10.

Example of Information in the Component Knowledge Base.

SYSTEM OPERATION

The system will be distributed with complete knowledge bases and the inference engine packaged to run on many of the personal computers already in wide use. The Policy knowledge base will be prepared by computer security experts in consultation with the policy-making organization. The initial implementation of this system will be for the DOE.

The system will use the knowledge bases to prompt the local security officer for information on the local computing environment. Once the information is collected the system will prepare a list of specific safeguards needed to secure the local computing system. The output will also contain, when available, the minimum, acceptable, and maximum implementation levels for each safeguard in the list.

All information is retained by the system to permit the security officer to explore possible changes in the local computing environment for cost-benefit studies or other changes in the list of safeguards.

FUTURE DIRECTIONS

Collecting the information for the various knowledge bases and developing the inference engine are the present efforts, but several significant possibilities for future extensions are already obvious.

The system could be easily used by an inspector or evaluator by simply entering the observed local computing environment and then using the output list of safeguards as a guide to confirm the presence or absence of the required safeguard.

Other possibilities are the automatic production of security plans and security test/certification plans.

SUMMARY

A knowledge-based system has been designed to collect and organize knowledge from computer security experts for use by a security officer. The system is policy-based and flexible to support changes in policy and technology.